Submitted to *Operations Research*
manuscript OPRE-2007-08-388

Fundamental Performance Limits for Multi-Stage Vehicle Routing Problems

Holly A. Waisanen, Devavrat Shah, and Munther A. Dahleh
Dept of EECS, MIT, waisanh@mit.edu, devavrat@mit.edu, dahleh@mit.edu

In the stochastic and dynamic multi-stage vehicle routing problem, a set of vehicles collectively transports and services demands which require service at multiple spatially separated locations. Arrival times and service locations of individual demands are not known until the demand's arrival, and vehicles are assumed to have infinite capacity and to travel with bounded velocity. The objective is to define real-time control policies for the vehicles to service the demands such that average demand time-in-system is minimized. Such problems arise in numerous applications, including courier services, inventory routing, and mobile sensor networks.

In this paper, we provide a general lower bound on the delay performance of a class of batching policies for the dynamic multi-stage vehicle routing problem. This requires us to overcome nontrivial technical difficulties presented by the non-work-conserving nature of the system and the service of multiple (possibly infinitely many) jobs simultaneously. The single stage Dynamic Traveling Repairperson Problem was first analyzed by Bertsimas and van Ryzin (1991). We recently examined the two-stage Dynamic Pickup and Delivery Problem in Waisanen et al. (2006). In the current paper, we extend the Dynamic Pickup and Delivery Problem to include a single inter-vehicle relay service between the pickup and delivery stages. For this four-stage problem, we present a policy for which the corresponding multi-stage lower bound is tight as a function of the scaling parameters (up to a constant). We note that our bounds are based on sample path analysis and hence are quite general.

Subject classifications: Networks/Graphs: stochastic, traveling salesperson, Probability: stochastic model applications, Transportation: vehicle routing
Area of review: Stochastic Models

1. Introduction

Multi-stage vehicle routing problems are characterized by a series of demands (messages) which require service at multiple locations within a given region by at least one of several mobile servers (vehicles). The vehicles are routed throughout the region to service and transport the demands such that some cost function is minimized. In the case that messages arrive with stochastic arrival times and service locations, a common cost function is the minimization of the expected time between a demand's arrival and its service completion over an infinite time horizon.

The single stage problem is a natural probabilistic formulation of the canonical vehicle routing problem, the Traveling Salesperson Problem (TSP). In the two stage problem, each demand must be picked up from one location and delivered to another, such as in an urban courier service. In this paper, we consider a four-stage extension of the pickup and delivery problem in which each demand is relayed once between a pickup and a delivery vehicle. In this single relay problem, each message requires four services: pickup, relay drop, relay pickup, and delivery. An example of such a problem is a courier service in which packages are dropped at a central sorting facility after pickup to be routed to the appropriate delivery vehicle.

Solutions to dynamic vehicle routing problems generally take one of two forms: nearest neighbor policies in which each vehicle choses the next demand to service based on the minimum distance from the current vehicel location, and batching policies in which demands are partitioned into groups to be collectively serviced according to the solution of the static vehicle routing problem formed by this subset of demands. Nearest neighbor policies are a natural analog to the Shortest Job First (SJF) policies for minimizing the delay in standard queueing systems. However, vehicle routing systems differ significantly from those studied in classical queuing theory.

In particular, many queuing control policies rely on knowledge of the actual realized service time that the server will spend on each message. Consider two ways of defining "service times" in the vehicle routing system. If service times are defined to be the time that the vehicle must spend onsite at the service location only, the system is not work conserving in the usual sense as vehicles are not performing work when traveling between demand locations. On the other hand, if service times are taken to be onsite service time plus the travel time from the current vehicle position, characterizing service time becomes intractable for any interesting policy as multiple demands are receiving service simultaneously while the vehicle is traveling. For the same reason, travel time may also be assigned to the service times of individual demands in other ways, but this assignment is far from trivial and nearest neighbor policies are difficult to analyze in general.

Batching policies, on the other hand, are simple to implement and analytically tractable. Typically, batches are formed by subsets of consecutive demands. Batches are serviced according to the solution of the static vehicle routing problem formed by this subset of demands. The analysis of a batching policy for a given vehicle routing problem provides an analytic upper bound on performance. However, until now, analytic lower bounds for batching policies have only been available in specific cases, most notably for the single stage problem.

The main contribution of this paper is a method of lower bounding the performance of batching policies for general multi-stage dynamic vehicle routing problems. This method takes into account the impact of both onsite service time and also travel time between demands. The key challenges of this analysis arise from the non-work-conserving nature of the system and and the simultaneous service of multiple (possibly infinitely many) jobs while the vehicle is traveling. This lower-bound method is applied to the single relay Dynamic Pickup and Delivery problem, but may readily be extended to more general multistage problems. Further, a batching policy is constructed to demonstrate that this bound is tight (up to a constant) for the single relay DPDP.

1.1. Related Work

The single-stage stochastic and dynamic vehicle routing problem is known as the Dynamic Traveling Repairperson Problem (DTRP), first analyzed by Bertsimas and van Ryzin (1991, 1993a,b). In the DTRP, each demand is serviced when one of the vehicles travels to the demand location and remains stationary at that point for a given onsite service time to "repair" the demand. Several extensions of the DTRP have been studied recently including policies for the DTRP with limited communication constraints in Frazzoli and Bullo (2004) and bounds on the DTRP with constraints on vehicle mobility in Savla et al. (2006).

Two-stage problems are commonly known as Pickup and Delivery Problems (PDP) in which each demand requires pickup service at its source location and then delivery service at its destination location elsewhere in region. PDPs arise in many practical applications including cab and courier services, manufacturing and inventory routing, less-than-truckload (LTL) trucking, emergency services, mobile sensor networks, and Unmanned Aerial Vehicle (UAV) routing. Recent survey papers on PDPs, both static and dynamic, include Savelsbergh and Sol (1995), Desaulniers et al. (2002), and Gendreau and Potvin (1998).

When demands are taken to be people, the PDP is also commonly referred to as the Dial-A-Ride Problem (DARP) (see survey Feuerstein and Stougie (2001)). Most research on the dynamic or online DARP (OLDARP) has used competitive analysis to quantify the loss in performance when information is not available a priori. In a significantly different approach, Swihart and Papastavrou (1999) and Waisanen et al. (2006) assume an infinite time horizon of continuously arriving messages and analyze the Dynamic Pickup and Delivery Problem (DPDP) directly, similar to the analytical results for the DTRP. While Swihart and Papastavrou (1999) examine the single vehicle DPDP, our earlier work in Waisanen et al. (2006) analyzes the multivehicle DPDP, providing lower bounds on the delay performance of the system and policies that achieve delay within a constant factor of these bounds. These bounds are somewhat problem specific in that they are valid only for a DPDP in which demands may not be relayed between vehicles. Here, the challenge is to develop a general framework for bounding the performance of any Dynamic Pickup and Delivery Problem without the relaying restriction. We will see that the single relay DPDP analyzed in the current paper improves average message delay by an order over that of the two-stage no-relay DPDP.

1.2. Organization

The remainder of this paper is organized as follows. Section 2 details the problem formulation. Section 3 presents the main result of the paper by establisting a lower bound for a class of batching policies. Section 4 describes and analyzes a single relay policy for the DPDP using a synchronous vehicle rendezvous schedule to relay the messages directly between vehicles. This establishes the tightness of our lower bound to within a constant factor. Finally, we discuss some further extensions in Section 5.

2. Problem Formulation

2.1. Model

Let there be n vehicles in a geographic area $\mathcal{A} = [0, \sqrt{A}]^2$. Each vehicle may move in any direction at any time with a velocity of magnitude $\leq v$. Messages are generated according to a stationary regenerative process with time intensity $\lambda(n)$. In particular, defining $\Lambda(t)$ to be the number of arrivals in the interval $[0,t]$ for a particular sample path ω, we have the following limit with probability 1:

$$\lim_{t\to\infty} \frac{\Lambda(t)}{t} = \lambda(n). \tag{1}$$

Each message j must receive service for a deterministic service time $\bar{s}(n)$ at each of at least two service locations: pickup at its source location and delivery at its destination. We assume that source and destination locations are fixed upon the arrival of the message and that these service locations are independent and have a uniform distribution on $\mathcal{A}$. Messages are serviced and transported between the service locations by the vehicles.

In addition to the pickup and delivery service, each message may also be transferred between vehicles via a set of relay services. We assume messages are relayed directly. That is, both vehicles involved in the relay must be colocated at an arbitrary service location for the full $\bar{s}(n)$ service time. Assume that for safety or other reasons, the colocated pair may perform the relay only if there are no other vehicle pairs within distance r of the relay service location. Because each vehicle provides an effective service, a direct relay counts as two services at a single location. Therefore, the single relay DPDP problem corresponds to a four-stage vehicle routing problem with the following stages: source pickup, relay drop, relay pickup, and destination delivery.

Note that $\bar{s}(n)$ is a fixed constant, but is expressed as a function of n to emphasize the connection between the arrival rate $\lambda(n)$, the number of servers n, and the maximum onsite service time that may be supported in a stable system. Further discussion of the stability condition may be found in Section 2.3.

Remark: We will use order notation to express the scaling of $\lambda(n)$ and $\bar{s}(n)$ as a function of n. Recall the following notation:

(i) $f(n)=O(g(n))$ means that $\exists$ a constant c and integer N such that $f(n) \leq cg(n), \forall n \geq N$.

(ii) $f(n)=\Omega(g(n))$ if $g(n)=O(f(n))$.

(iii) $f(n)=\Theta(g(n))$ means that $f(n)=O(g(n))$ and $g(n)=O(f(n))$.

For ease of notation, at various places in the paper we will use λ for $\lambda(n)$.

2.2. Control Policies

We define a control policy to be a set of decision making rules that determines the pickup, relaying, and delivery schedule of messages as they arrive, based on a set of constraints on the information available to each of the vehicles and the system as a whole. The lower bounds presented in this paper are independent of information constraints. In Section 4, we will demonstrate that these bounds are tight for the four-stage single relay DPDP even when inter-vehicle information exchange is prohibited after a centralized initialization period.

We limit our attention to *batching* policies.

Definition 1 (Batch). A batch is a set of requests for service, such that 1) all service requests within a single batch are assigned to a single vehicle, and 2) once a vehicle begins service of one of the requests in the batch, it completely serves all the requests in the batch, oblivious to other demands in the system.

Definition 2 (Batching Policy). Under a batching policy, each request for service is buffered at a batch processor upon arrival. Service requests are assigned to batches in some arbitrary way (with a few assumptions listed below), and a request remains at the batch processor until the batch it is assigned to is released into the batch queue. Once a batch is released to the batch queue, no new service requests may be added or removed from the batch. Vehicles serve the batches from the batch queue one at a time.

Batches are numbered in order of the release of the batch to the batch queue. Let B_k be the number of requests contained in the k^{th} batch. The service of each batch has two components: 1) the service of the B_k individual messages, and 2) overhead time required to complete the batch service. Denote this overhead time by I_k. The total time to service the k^{th} batch is denoted as T_k, which is a function of B_k and I_k.

Assume that messages are assigned to batches in such a way that the time averages of B_k, B_k^2, B_kI_k, and I_k exist and are finite with probability 1, i.e.

$$\overline{B} \triangleq \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} B_k < \infty, \quad w.p.1$$
$$\overline{B^2} \triangleq \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} B_k^2 < \infty, \quad w.p.1$$
$$\overline{BI} \triangleq \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{K} B_kI_k < \infty, \quad w.p.1$$

$$\overline{I} \triangleq \lim_{K\to\infty} \frac{1}{K}\sum_{k=1}^{k} I_k < \infty, \quad w.p.1.$$

For technical reasons, we will require the number in the batch and the idle time to be related by

$$\frac{\overline{B^2}}{\overline{B}} \geq \frac{\overline{BI}}{\overline{I}}.$$

Specifically, for the single relay DPDP, each message generates two requests for service which arrive to the system at different times: the pickup request issued upon arrival and the delivery request issued upon message relay. Therefore, the set of arrived requests include both external arrivals and also internal arrivals of demands as they are relayed between vehicles. The overhead, I_k, is defined to be the traveling time between the service locations of all the messages in the batch. Each message requires exactly two services from each vehicle that it is assigned to (either pickup/relay or relay/delivery). Therefore, the total time to service batch k is

$$T_k = 2\bar{s}(n)B_k + I_k.$$

2.3. Performance Metrics

There are two main performance measures to be defined: stability and average delay. Informally, a system is *stable* if the messages arriving to the system have finite average delay. The precise average delay differentiates the performance between various stable policies. Before precisely defining average delay and stability, we introduce some preliminary definitions.

Message j is said to be *assigned* to vehicle i at time t if the message is waiting for pickup by vehicle i or has already been picked up by vehicle i but has not yet been completely serviced by that vehicle.

$$\mathbf{1}_{j,i}(t) = \begin{cases} 1 & \text{if message } j \text{ is assigned to vehicle } i \text{ at time } t \\ 0 & \text{else} \end{cases}$$

A vehicle is *traveling* if it is moving between service locations. A vehicle is in *onsite* service when it is stopped at a service location and performing pickup, relay or delivery service.

$$\mathbf{1}_{i,T}(t) = \begin{cases} 1 & \text{if vehicle } i \text{ is traveling at time } t \\ 0 & \text{else} \end{cases}$$

$$\mathbf{1}_{i,O}(t) = \begin{cases} 1 & \text{if vehicle } i \text{ is in onsite service at time } t \\ 0 & \text{else} \end{cases}$$

We assume that at any time there are messages in the system, the vehicle is either traveling or in onsite service, i.e.

$$\mathbf{1}_{j,i}(t) = \mathbf{1}_{j,i}(t)\left[\mathbf{1}_{i,T}(t) + \mathbf{1}_{i,O}(t)\right].$$

These indicator functions may be used to define various measures of delay for an individual message j, along with their limiting expectation.

The delay of message j at vehicle i while i is traveling is

$$W_{T,i}(j) = \int_0^\infty \mathbf{1}_{j,i}(\tau)\mathbf{1}_{i,T}(\tau)d\tau$$

Likewise, the delay of message j at vehicle i while i is in onsite service is

$$W_{O,i}(j) = \int_0^\infty \mathbf{1}_{j,i}(\tau)\mathbf{1}_{i,O}(\tau)d\tau$$

This onsite service time includes both the two services associated with the message itself as well as any other services that occur while j is assigned to vehicle i.

Therefore, the delay of a single message j while it is either in service or in queue for vehicle i is

$$W_i(j) = W_{T,i}(j) + W_{O,i}(j).$$

The total delay of a single message j is then

$$\begin{aligned} W(j) &= \sum_{i=1}^{n} W_i(j) \\ &= W_T(j) + W_O(j) \end{aligned}$$

where $W_T(j) = \sum_{i=1}^{n} W_{T,i}(j)$ and $W_O(j) = \sum_{i=1}^{n} W_{O,i}(j)$.

The average delay over all messages that pass through the system is then

$$W = \limsup_{J\to\infty} \frac{\sum_{j=1}^{J} W(j)}{J}.$$

Define the subsequence of messages served by vehicle i as:

$$X_i = \{(i_1, i_2, ..., i_j, ...) : W_i(i_j) > 0\}.$$

Then the total delay at vehicle i for messages served by vehicle i is:

$$W_i = \limsup_{J\to\infty} \frac{\sum_{j=1}^{J} W(i_j)}{J}$$

Definition 3 (Stability). If the limits above exist, and are finite, the system is *stable* and the total delay W is composed of two parts:

$$W = W_T + W_O$$

where

$$\begin{aligned} W_T &= \lim_{J\to\infty} \frac{\sum_{j=1}^{J} W_T(j)}{J} \\ W_O &= \lim_{J\to\infty} \frac{\sum_{j=1}^{J} W_O(j)}{J}. \end{aligned}$$

Let $\Lambda_i(t)$ be the number of arrivals to vehicle i in the interval $[0,t]$, including both external arrivals to the system at vehicle i (new pickups) and also messages that are relayed to vehicle i for further service. It is possible for a message to be relayed from one vehicle to itself if the same vehicle handles both the pickup and the delivery of the message. For ease of exposition, this relay is counted as a new internal arrival the vehicle when the message is picked up. Define the time average rate of arrivals to vehicle i to be

$$\lambda_i = \lim_{t\to\infty} \frac{\Lambda_i(t)}{t} \tag{2}$$

where this limit is assumed to exist. Because relays increase the number of times each message is served,

$$\sum_{i=1}^{n} \lambda_i \geq \lambda.$$

For the single relay system in which each message is handled by exactly two vehicles

$$\sum_{i=1}^{n} \lambda_i = 2\lambda.$$

Assume that W_i and $W_{i,O}$ are both increasing functions of λ_i. That is, when a vehicle serves proportionally more messages, the average delay seen by messages served by that vehicle also increases. This is a natural assumption in the case where service locations are uniformly distributed and all onsite service times are iid.

With these limits and the definitions of λ and λ_i in (1) and (2) respectively, we have the following equivalent representation for W.

$$\begin{aligned} W &= \lim_{J\to\infty} \frac{\sum_{j=1}^{J} W(j)}{J} = \lim_{t\to\infty} \frac{\sum_{j=1}^{\Lambda(t)} W(j)}{\Lambda(t)} \\ &= \lim_{t\to\infty} \sum_{i=1}^{n} \frac{\Lambda_i(t)}{\Lambda(t)} \frac{\sum_{j=1}^{\Lambda_i(t)} W(i_j)}{\Lambda_i(t)} \\ &= \lim_{t\to\infty} \sum_{i=1}^{n} \frac{\Lambda_i(t)}{t} \frac{t}{\Lambda(t)} \frac{\sum_{j=1}^{\Lambda_i(t)} W(i_j)}{\Lambda_i(t)} \\ &= \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_i \end{aligned} \tag{3}$$

Similar expressions may also be obtained for W_O and W_T.

$$\begin{aligned} W_O &= \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_{O,i} \\ W_T &= \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_{T,i} \end{aligned} \tag{4}$$

Viewing the network of vehicles as a non-work-conserving, n-server system with service times defined to be onsite service only, define the server and system utilizations as

$$\begin{aligned} \rho_i &= \lambda_i \bar{s} \\ \rho &= \frac{\sum_{i=1}^{n} \lambda_i \bar{s}}{n} = \frac{\sum_{i=1}^{n} \rho_i}{n} \end{aligned} \tag{5}$$

Note that while ρ is a function of the system parameters and the number of relays per message, the individual ρ_i are also a function of the specific control policy in effect. A necessary condition for a policy to be stable in this setup is

$$\rho_i \leq 1, \forall i \tag{6}$$

This necessary stability condition is derived by comparing each vehicle to a work-conserving $M/D/1$ queue with arrival rate λ_i and service time $2\bar{s}(n)$.

Further, we demonstrate by construction in Section 4 that the following is a sufficient condition for the existence of a stable policy for the single relay DPDP:

$$\rho \leq \frac{4\lambda(n)\bar{s}(n)}{n} < 1. \tag{7}$$

If each message is served by exactly two vehicles (all messages are relayed), then the left hand inequality of (7) holds with equality.

As noted in the analysis of the DTRP in Bertsimas and van Ryzin (1991), the stability condition does not depend on the geometry of the system, i.e. the placement of the message service locations, but only on the net arrival rate of onsite workload. This stability condition extends for our case as well.

2.4. Problem Statement

We will call the above defined control problem the single-relay Dynamic Pickup and Delivery Problem (DPDP). The goal is to compute a tight lower bound on the average message delay, W, under any stabilizing batching control policy for the single-relay DPDP for all ranges of the scaling parameters, $\lambda(n), n, A$, and $\bar{s}(n)$.

The tightness of the lower bounds is demonstrated by the construction of a valid control policy for each vehicle that decides the pickup, relay and delivery schedule of arriving messages such that each message is served by at most two vehicles, one for pickup and one for delivery, and the average message delay is of the same order as that of the lower bound.

3. Lower Bounds on Delay Performance

In this section we prove a lower bound on the expected delay for the single-relay Dynamic Pickup and Delivery Problem. This lower bound comprises three terms, each of which independently provides a lower bound the delay for all scaling parameters. However, each bound correctly characterizes the qualitative behavior of the delay only over a subset of the scaling parameters. The key contribution here is the the combination of these bounds to provide a tight lower bound over a wider range of scaling regimes.

To review, we have the following set of assumptions:

(P1) The arrival process is a stationary regenerative process with rate $\lambda(n)$.

(P2) Service locations are uniform and independently distributed. Service times are deterministic of length $\bar{s}(n)$.

(P3) λ_i and W_i are positively correlated, that is, $\lambda_i \geq \lambda_{i'} \implies \overline{W}_i \geq \overline{W}_{i'}$

(P4) A batching policy is used in which the sample path average number in batch and batch overhead time are related as follows:

$$\frac{\overline{B^2}}{\overline{B}} \geq \frac{\overline{BI}}{\overline{I}}, \tag{8}$$

(P5) The stability condition $\rho = 4\lambda(n)\bar{s}(n)/n < 1$ holds.

We also introduce the following two additional assumptions:

(P6) The arrival process is Poisson with rate $\lambda(n)$.

(P7) The arrival rate and number of servers scale as $\lambda(n) = \Omega(n)$.

We are now ready to state the main theorem as follows:

THEOREM 1. *For any four-stage Dynamic Pickup and Delivery system satisfying properties (P1)-(P5), the average system delay is bounded from below as*

$$W = \Omega\left(\frac{\sqrt{A}}{v(1-\rho)}\right) + \Omega(\bar{s}(n)). \tag{9}$$

If properties (P6)-(P7) also hold, then the following is a tight lower bound

$$W = \Omega\left(\frac{\lambda(n)A}{v^2(1-\rho)^2n^2}\right) + \Omega(\bar{s}(n)) + \Omega\left(\frac{\sqrt{A}}{v(1-\rho)}\right). \tag{10}$$

The first part of the proof is a new lower bound on the DPDP, which is combined with a trivial lower bound (the service time) to derive (9). This proof method may be extended to more general multi-stage vehicle routing problems. The second part of the proof yields the first term of equation (10) and is just a restatement of the DTRP result of Bertsimas and van Ryzin (1993a).

3.1. Universal Lower Bound

First note that the total service time of a message, $4\bar{s}(n)$, is a lower bound on the time a message spends in the system, that is

$$W = \Omega(\bar{s}(n)). \tag{11}$$

This lower bound may be combined in a convex combination with the following lower bound to derive (9).

THEOREM 2. *For any stable batching policy for the single relay DPDP for which properties (P1)-(P5) hold, the average system delay is bounded as*

$$W = \Omega\left(\frac{\sqrt{A}}{v(1-\rho)}\right).$$

The proof of Theorem 2 rests on the analysis of a reduced batch system which is coupled to the original DPDP system in that delay in the reduced system is a lower bound on the delay in the DPDP system.

DEFINITION 4 (REDUCED BATCH SYSTEM). The reduced batch system is a system under a batching policy in which messages at the batch processor or in the batch queue are ignored. That is, in the reduced system, service requests do not arrive until just before the batch to which they belong begins service. Comparing this to the original system, delay between the time a message arrives and is assigned to a vehicle and the beginning of its batch service is ignored. Because two intervals of delay are ignored for each message, one for each service request generated, the delay of the reduced system provides a lower bound on the delay of the overall DPDP system.

The following proof comprises an analysis of the reduced system.

Before beginning the proof of Theorem 2, suppose we have the following Lemma bounding the time spent in onsite service:

LEMMA 1. *If inequality* (8) *holds for a batching policy for the single relay DPDP, then for the reduced system*

$$W_O \geq \frac{\rho}{2-\rho} W$$

with $\rho = 4\lambda(n)\bar{s}(n)/n < 1$.

Note that this is trivially satisfied for a work conserving system in which the vehicle is in service anytime there are active messages in the system, that is, $W_O = W$. It is the travel times between service locations that make the DPDP a non-work conserving system.

We now begin the proof of Theorem 2.

[Proof of Theorem 2] Recall the following definitions: B_k and I_k are the number of messages and the travel time associated with batch B_k. Because a vehicle is stationary whenever it is performing onsite service, message delay may be partitioned into time spent traveling plus time spent in onsite service, i.e. $W = W_T + W_O$.

The average total travel time, W_T, is lower bounded by the expected straight line distance between the message's source and destination. Because sources and destinations are independently and uniformly distributed, this distance is $c_1\sqrt{A}$, where the constant $c_1 \approx 0.52$ (see Larson and Odoni (1981), p. 135).

Combining this expected travel time with the sample path average of Lemma 1 we have

$$\begin{aligned} W &\geq W_T + W_O \\ &\geq c_1\sqrt{A}/v + \frac{\rho}{2-\rho} W \\ &= \frac{2-\rho}{2} \frac{c_1\sqrt{A}}{v(1-\rho)} \\ &\geq \frac{1}{2} \frac{c_1\sqrt{A}}{v(1-\rho)} \end{aligned}$$

where $\rho = \frac{4\lambda(n)s(n)}{n}$. The final equation is the desired result of Theorem 2.

Before proving Lemma 1, we present and prove Lemma 2 below which makes a statement about individual vehicles. The proof of Lemma 1 will use these single vehicle bounds to make a statement about the overall system delay. First, define $N_{O,i}$ to be the time average number in the vehicle queue when the vehicle is in onsite service and $N_{I,i}$ to be the time average number when the vehicle is not in service of an individual message (either traveling or idling). Assume that both averages exist and are finite.

LEMMA 2. *The following inequality holds for the reduced system under any batching policy which satisfies* (8)*:*

$$N_{O,i} \geq \frac{N_{I,i}}{2} \tag{12}$$

for all vehicles i in the multi-vehicle system. That is, the average number assigned to a vehicle when the vehicle is not in service is no more than twice the average number when the vehicle is in service.

[Lemma 2] Examine a single batch k_i that is served by vehicle i with $B_{k,i}$ messages and $I_{k,i}$ total travel time between service locations to fully serve the batch. For ease of notation, we will drop the references to the vehicle i for the remainder of the proof of this Lemma.

By the definition of the batching policy, all of the B_k messages must be assigned to the vehicle before any deliveries of these messages can occur. Because subsequent arrivals are ignored in the reduced system, this implies that the number in the reduced system is exactly B_k at the beginning of the batch and at most B_k for the duration of the batch. Therefore, letting $N_{I,k}$ be the number of messages in the system during the travel time in this batch, $N_{I,k} \leq B_k$.

Further, during the delivery phase, the number in the system falls from exactly B_k to 0 in unit decrements. The decrement from value $B_k - j + 1$ to $B_k - j$ occurs immediately after $\bar{s}(n)$ time has been spent serving the jth message. Averaging over each of these increments, the average value is exactly half of the peak value, that is, $N_{O,k} = B_k/2$. Therefore, the inequality (12) holds over each of the batches individually.

Now compute the expectation of N_O and N_I over the first K batches, where K is some arbitrary large number. Because $2\bar{s}(n)$ service time is required for each message, the total time in onsite service is exactly $N_{O,k} 2\bar{s}(n)$. Therefore, each batch in N_O sum is weighted by the number of messages served by that batch, B_k.

$$\begin{aligned} N_O &\geq \sum_{k=1}^{K} \frac{B_k}{2} \frac{B_k}{\sum_{k=1}^{K} B_k} \\ &= \sum_{k=1}^{K} \frac{B_k^2}{2\sum_{k=1}^{K} B_k} \quad \overset{K\to\infty}{\longrightarrow} \quad \frac{\overline{B^2}}{2\overline{B}} \end{aligned} \tag{13}$$

On the other hand, in N_I, the batches are weighted by the time spent traveling. N_I also includes idle time between batches that is not accounted for in the sums below. Accounting for this idle time would only further reduce the actual value of N_I.

$$N_I \leq \sum_{k=1}^{K} B_k \frac{I_k}{\sum_{k=1}^{K} I_k} \quad \overset{K\to\infty}{\longrightarrow} \quad \frac{\overline{BI}}{\overline{I}} \tag{14}$$

Combining equations (13) and (14), and recalling the inequality (8), we have

$$N_O = \frac{1}{2}\frac{\overline{B^2}}{\overline{B}} \geq \frac{\overline{BI}}{\overline{I}} \geq \frac{N_I}{2}.$$

Therefore, we have proven that when (8) holds, (12) holds as well.

Finally, we prove Lemma 1.

[Lemma 1] Define $T_{O,i}$ to be the total time in the interval $[0, T]$ in which the vehicle i is in service.

$$\begin{aligned} N_i &\leq \lim_{T\to\infty} \frac{\int_0^T N_i(t)\left[\mathbf{1}_{O,i}(t) + \mathbf{1}_{I,i}(t)\right] dt}{T} \\ &= \lim_{T\to\infty} \frac{\int_0^T N_i(t)\mathbf{1}_{O,i}(t)dt}{T_{O,i}} \left(\frac{T_{O,i}}{T}\right) + \frac{\int_0^T N_i(t)\mathbf{1}_{I,i}(t)dt}{T - T_{O,i}} \left(\frac{T - T_{O,i}}{T}\right) \\ &\leq N_{O,i} \lim_{T\to\infty} \frac{T_{O,i}}{T} + N_{I,i} \lim_{T\to\infty} \frac{T - T_{O,i}}{T} \end{aligned}$$

Stability implies that

$$\begin{aligned} \lim_{T\to\infty} \frac{T_{O,i}}{T} &= \lim_{T\to\infty} \frac{\Lambda_i(T) 2\bar{s}(n)}{T} \\ &= 2\lambda_i \bar{s}(n) = \rho_i \end{aligned}$$

Therefore N_i, the unconditioned average number assigned to vehicle i, is given by

$$N_i \leq \rho_i N_{O,i} + (1 - \rho_i) N_{I,i}.$$

Substituting in the result of Lemma 2 above yields

$$\begin{aligned} N_i &\le \rho_i N_{O,i} + 2(1-\rho_i)N_{O,i} \\ &\le (2-\rho_i)N_{O,i} \end{aligned}$$

and therefore

$$N_{O,i} \ge \frac{N_i}{2-\rho_i}. \tag{15}$$

Given the definition of λ_i in (2) and the assumption of the existence of the sample path average N_i, Little's Law may be applied to the single vehicle system, $N_i = \lambda W_i$. Little's Law may also be applied to the onsite system formed by deleting all times in which the system is not in service. Because messages are always being served in the onsite system, messages complete service in the onsite system at a fixed rate of $1/\bar{s}$ with no idling. Because half of these service completions are departures, this implies that the departure rate from the onsite system is $1/2\bar{s}$. For stability, the arrival rate to the onsite system must also be $1/2\bar{s}$, and the corresponding Little's Law result is $N_{O,i} = W_{O,i}/2\bar{s}$. See Appendix A for the full details of this proof. Therefore, the following is equivalent to equation (15):

$$\begin{aligned} \frac{1}{2\bar{s}} W_{O,i} &\ge \frac{\lambda_i W_i}{2-\rho_i} \\ W_{O,i} &\ge \frac{\rho_i}{2-\rho_i} W_i \end{aligned}$$

for all vehicles i.

For each vehicle, this may then be rearranged as

$$W_{O,i} \ge \frac{1}{2}\rho_i \left(W_i + W_{O,i}\right)$$

Then, taking the weighted sum of these terms and applying the definitions in equations (3), (4), and (5),

$$\begin{aligned} \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_{O,i} &\ge \frac{1}{2} \sum_{i=1}^{n} \left[\frac{\lambda_i}{\lambda} \rho_i \left(W_i + W_{O,i}\right)\right] \qquad (16) \\ W_O &\ge \frac{1}{2} \frac{\sum_{i=1}^{n} \rho_i}{n} \left[\sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_i + \sum_{i=1}^{n} \frac{\lambda_i}{\lambda} W_{O,i}\right] \qquad (17) \\ &\ge \frac{1}{2}\rho(W + W_O) \\ &\ge \frac{\rho}{2-\rho} W \end{aligned}$$

The implication of (17) from (16) is given by the assumption that W_i and $W_{i,O}$ are both increasing functions of λ_i (and therefore of ρ_i), and noting that this positive correlation implies that

$$\frac{1}{n} \sum_{i=1}^{n} \rho_i \sum_{i=1}^{n} \lambda_i W_i \le \sum_{i=1}^{n} \rho_i \lambda_i W_i$$

and likewise for $W_{i,O}$.

Therefore, equation (12) is sufficient to show the inequality for the multi-vehicle system and complete the proof of Lemma 1.

The following Corollary provides some more intuitive conditions under which (8) and Theorem 2 hold.

COROLLARY 1.

$$W = \Omega\left(\frac{\sqrt{A}}{v(1-\rho)}\right)$$

for any batching policy such that $\rho = \frac{4\lambda\bar{s}}{n} < 1$*, the time averages of* $E[B_k], E[B_k^2], E[B_kI_k]$*, and* $E[I_k]$ *exist, and at least one of the following conditions is satisfied:*

(a) I_k *and* B_k *are uncorrelated or negatively correlated random variables.*

(b) $E[B_k|I_k]$ *is either a constant or a linear function of* I_k*, i.e.* $E[B_k|I_k] = \delta$ *or* $E[B_k|I_k] = \gamma I_k$ *for some* $\delta, \gamma \in [0, \infty)$*.*

(c) $E[I_k|B_k]$ *is an affine function of* B_k^α *for some* $\alpha \in [0, 1]$*, i.e.* $E[I_k|B_k] = \gamma B_k^\alpha + \delta$ *for some* $\delta, \gamma \in [0, \infty), \alpha \in [0, 1]$*.*

The time average conditions are required to relate the expected values of the batch parameters to the time averages required for Theorem 2. Intuitively, condition **(c)** in Corollary 1 means that the overhead time per batch does not grow faster than linearly with the number in batch. This is a natural condition for batch service as it implies that there is some economy of scale associated with grouping demands into batches for service.

The proof of Corollary 1 may be found in the appendix.

3.2. Lower Bound from the Dynamic Traveling Repairperson Problem (DTRP)

Because the DTRP is a single stage problem, equivalent to message pickup only, the DTRP delay performance is naturally a lower bound on the performance of any multi-stage problem.

When properties (P2) and (P6) hold, and $\rho = \lambda(n)\bar{s}(n)/n < 1$ the delay of the DTRP is lower bounded in Bertsimas and van Ryzin (1993a), Theorem 2 as

$$W_{DTRP} = \Omega\left(\frac{\lambda A}{v^2(1-\rho)^2n^2}\right) - \Omega\left(\frac{\bar{s}(1-2\rho)}{\rho}\right) \tag{18}$$

Apart from the necessary change in the stability condition from (P5) to $\rho = \lambda(n)\bar{s}(n)/n < 1$, the first term of this lower bound matches the first term in the delay for the multistage DPDP in equation (10).

Further, if property (P7) holds, the second term of (18) becomes

$$\begin{aligned}
-\Omega\left(\frac{\bar{s}(1-2\rho)}{\rho}\right) &= \Omega\left(2\bar{s} - \frac{\bar{s}}{\rho}\right) \\
&= \Omega\left(2\bar{s} - \frac{n}{\lambda}\right) \\
&= \Omega(\bar{s})
\end{aligned}$$

This matches the second term of (10).

Finally, these two terms from the DTRP may be combined in convex combination with (2) to complete the derivation of (10).

3.3. Comparison of Lower Bounds

Examining the first two terms of the lower bound in Theorem 1, note that for $\lambda(n) = O\left(n^2(1-\rho)^2\frac{v^2}{A}\right)$ the lower bound provided by the DTRP goes to $\Omega(\bar{s}(n))$ as n becomes large. Because each message must be carried at least from source to destination, zero travel delay is not feasible and this bound is not tight. When $\lambda(n) = \Omega\left(n^2(1-\rho)\frac{v}{\sqrt{A}}\right)$, the universal lower bound for batching policies is dominated by the DTRP lower bound, and therefore this bound is not tight over all parameter scalings either. For small arrival rates with moderate ρ, the onsite service time may dominate the other two terms. All three bounds are required to provide a qualitatively meaningful lower bound over the wider range of parameter scalings.

Specifically, suppose that λ/n^2 scales as $\Theta\left((1-\rho)^\delta\right)$ for some $\delta \in \mathbb{R}$, and assume that $\sqrt{A}/v = \Theta(1)$. For $\delta > 1$, the first term of the lower bound in Theorem 1 is dominated by the third term. Since this holds without the scaling requirement of $(P7)$, the lower bound is tight for all scalings with $\lambda/n^2 = O\left((1-\rho)^1\right)$. For scalings in which $\lambda/n^2 = \Omega\left((1-\rho)^1\right)$, the lower bound is tight only if $\lambda/n = \Omega(1)$ as well.

Note that the delay W in the DTRP result in (18) refers to the expected delay of a randomly selected message. By the law of large numbers, if this expected value exists, then it is equal to $W = \lim_{J\to\infty}\sum_{j=1}^{J} W_j/J$ as defined in this paper with probability 1. On the other hand, the universal lower bound in (2) is proven directly on a sample path basis for any system realization ω. Assuming the appropriate regenerative properties, the lower bound in (2) holds when time averages are replaced by expectations as well.

4. Upper Bounds for Single-Relay Policies

In this section, we demonstrate the tightness of the lower bound of Theorem 1 for the four-stage single relay DPDP. We present and analyze a policy which uses a synchronous vehicle rendezvous schedule to relay messages directly between vehicles. Under certain assumptions on the vehicle rendezvous locations, this policy achieves delay with order that is arbitrarily close to that of the lower bound.

4.1. Policy Description

This policy has two general components, Assignment and Service, with the Service component being carried out in three phases.

4.1.1. Assignment A spatially based assignment policy is used to allow arriving messages to be assigned to vehicles without any real-time communication between the vehicles. This assignment policy may be initialized by a centralized controller and then implemented in a decentralized manner by each of the vehicles.

The region is divided into a $\sqrt{n} \times \sqrt{n}$ grid of cells, each of area A/n. Exactly one vehicle is assigned to each cell and is responsible for performing all of the pickups and deliveries in that cell. Upon arrival from outside the system, a message is assigned to the vehicle responsible for the cell in which the message's source location lies. When a message is relayed from the pickup to the delivery vehicle, it is immediately assigned to the vehicle responsible for the cell containing the message's destination.

4.1.2. Service Each vehicle has the same basic service policy, differing only by assignment region. The following service policy is described for a single vehicle. This policy for each vehicle has three basic steps: Pickup Batching and Service, Relay and Delivery Batching, and Delivery Service. Each vehicle cycles through these steps as long as there are messages to be served. In cycle k, two kinds of batches are defined: B_k, the pickup batch, and B_k^*, the delivery batch.

Pickup Batching and Service: In order to maintain a synchronous vehicle rendezvous schedule, messages are batched in such a way that the total time to service each batch may be deterministically upper bounded. For this, each vehicle maintains n source-destination queues of messages, one for each of the n cells in which the destination locations of arriving messages may occur. The k^{th} pickup batch, B_k, is formed by collecting up to the first N_n messages from each of the n queues to form a batch, where N_n is a parameter to be determined. The total number of messages to be serviced in each pickup batch is then at most $B_k \leq N_T = nN_n$.

To service a batch of messages, the vehicle computes and then traverses a worst case Traveling Salesperson (TSP) tour through the source locations of all of the messages contained in the batch, pausing at each service location to pickup the corresponding message. A worst case tour is used to maintain the deterministic synchronicity between the vehicles.

Relay and Delivery Batching: To relay messages to their delivery vehicles, a pre-determined synchronous schedule is used such that each vehicle meets up with every other vehicle during each batch service time to hand off the appropriate messages (see Lemma 3). The rendezvous points at which the vehicles meet are predetermined and are distributed throughout the region. Assume that for safety or other reasons, the minimum interpoint distance between valid rendezvous points is r.

As the messages are being relayed, the delivery batch is collected by receiving at most N_n messages from each of the other vehicles. Therefore, the number in the delivery batch is deterministically bounded by $B_k^* \leq N_T = nN_n$ as well.

Delivery Service: Once the inter-vehicle meetings are complete and all messages to be delivered have been received, another worst-case TSP tour is performed through the destination locations of the messages in the pickup batch.

These three stages of batch service occur within constant time-length $T_k = T$. N_n and T are the policy parameters to be determined.

4.2. Performance Analysis

THEOREM 3. *For the No-Depot policy described above, the delay scales as*

$$W = O\left(\frac{\lambda(n)A}{v^2(1-\rho)^2n^2}\right) + O\left(\frac{\sqrt{A}+nr}{v(1-\rho)}\right).$$

[Theorem 3] To establish validity of the above described policy as well as analyze its performance, we need the following two Lemmas. The first lemma establishes existence of a synchronous schedule for rendevous between vehicles and the second lemma establishes a worst case bound on TSP tour.

LEMMA 3. *Given N locations arbitrarily located in a square region of area B, there exists a tour through these points of length at most $2\sqrt{2NB}$.*

[Lemma 3] First note that if $N = 1$, the total time to visit the location and then return to the starting point, starting from anywhere in the region, is at most $2\sqrt{2B}$, so the bound in the theorem holds. The following is for $N \geq 2$.

Divide the region into N cells of area B/N. Consider a tour that begins at an arbitary location, then travels directly to the center of the upper-leftmost cell. The tour then travels between the centers of all the cells in a row-by-row manner, working across and then down through the region. Once all cell centers have been visited, the vehicle returns to the starting point to complete the tour. Such a tour through the cell centers takes at most time $N\sqrt{\frac{B}{N}} + \sqrt{2B} = \sqrt{(N+2)B} \leq \sqrt{2NB}$ for $N \geq 2$.

The tour through the N arbitrarily located points is performed by following the cell tour above, but stopping in each cell to visit all of the required locations that are located within that cell. To visit each location, the vehicle travels from the cell center to the the location and then back to the cell center. Each of these visits takes at most $\sqrt{2B/N}$. Since there are N locations to visit in this way, the location visits take a total of at most $\sqrt{2BN}$ in addition to the cell tour.

Combining this with the cell tour length above, the total tour through the N locations takes at most $2\sqrt{2NB}$.

Lemma 4. *Given n vehicles, there exists a schedule of length n such that each vehicle visits all other $n-1$ vehicles at least once.*

[Lemma 4] Consider a complete bipartite graph of $2n$ nodes, where each vehicle is represented by one node on the left and one on the right. An edge between node i on left and node j on right represents the requirement that vehicle i must meet vehicle j.

Now color the edges of this graph such that no two edges connected to the same node have the same color. By assigning color k to the edge between vehicle i on the left and vehicle $(i+k)$ mod n on the right, this may be accomplished using n colors. The schedule is then constructed by letting each color represent a time slot in which the two vehicles are assigned to meet and transfer messages.

Next, we use these lemmas to obtain appropriate values of T, N_T so that all the arriving messages are eventually delivered to their destinations and we will evaluate the induced delay. Given any $\varepsilon > 0$, let

$$N_T = \frac{(1+\varepsilon)\lambda(n)}{n} T.$$

Let T_{TSP} be the worst-case travel time it takes to tour-through pickup or delivery locations of N_T messages in cell of area A/n. Then by Lemma 3

$$T_{TSP} \leq \sqrt{\frac{8N_T A}{n}} \; = \; \sqrt{\frac{8(1+\varepsilon)\lambda(n)A}{n^2}}.$$

To complete the description of the rendezvous schedule guaranteed by Lemma 4, the locations of the vehicle rendezvous must be specified. If there is no restriction on the separation of rendezvous points, all vehicles may be assumed to travel to the center of the region and perform the handoffs without any need for further travel. If the pairs of vehicles must be separated by at least r, then the time for each vehicle to visit each of the rendezvous points is at least nr/v. If $r = 0$, all vehicles may perform the rendezvous at the center of the region for a total travel time upper-bounded by $\sqrt{2A}/v$. If the rendezvous point for each vehicle pairing represented may be taken to the be the center of the cell assigned to the vehicle on the left hand side in the graph, then the distance between points is at most $\sqrt{2A}$ and the total travel may be upper-bounded by $n\sqrt{2A}/v$, or $r = \sqrt{2A}$. Depending on the precise requirements of the system, assume that we may choose a set of relay locations separated by at least r such that $(\sqrt{2A} + nr)/v$ is an upper bound on the trave time through these locations.

The total time to exchange the messages during the rendevous of vehicles (both relaying and receiving) is $2N_T\bar{s}(n)$. Hence, the total batch time T can be bounded above as

$$\begin{aligned} T &\leq 2T_{TSP} + 4N_T\bar{s}(n) + \frac{\sqrt{2A}+nr}{v} \\ &= \sqrt{\frac{32(1+\varepsilon)\lambda(n)A}{n^2}} + (1+\varepsilon)\rho T + \frac{\sqrt{2A}+nr}{v}. \end{aligned} \tag{19}$$

From (19) and some manipulation will lead to the conclusion that it is sufficient to have T such that

$$T = O\left(\frac{\lambda(n)A}{(1-(1+\varepsilon)\rho)^2 n^2}\right) + O\left(\frac{\sqrt{A}+nr}{(1-(1+\varepsilon)\rho)}\right). \tag{20}$$

To complete the proof, we examine each of the arrival queues, show that they are stable with N_n as given, and then compute the time a message spends waiting to be collected into a batch. Note that in time T, in a given cell $\lambda(n)T/n^2$ messages arrive that are destined for any other cell. In the above described scheme with the selection of T as in (20), each vehicle serves up to $N_n = (1+\varepsilon)\lambda(n)T/n^2$ messages for a given pair of cells. Thus, we have a service rate higher than the arrival rate and hence by standard queueing argument, it must be stable.

Finally, we compute the average delay per message in this scheme. To this end, note that each message has the following types of delays: (a) waiting time to be serviced in a cell after arrival and (b) the batch time T. Now, the T is bounded above as (20). To bound (a), note that messages are queued separately depending on their destination cells. Consider one particular queue for a destination cell. The arrivals to this queue happen at rate $\lambda(n)/n^2$ while every T units of time, $N_n = N_T/n$ of them get served. Delay through this queue can be upper bounded by T plus the delay through an M/D/1 queue with arrival rate $\lambda(n)/n^2$ and deterministic service requirement of $T/N_n = \frac{n^2}{(1+\varepsilon)\lambda(n)}$.

The Pollacek-Khinchin formula in Wolff (1989) may be applied to compute the average delay in this M/D/1 queue to be $O(1/\varepsilon)$. Therefore the total average delay experienced by message between arrival and delivery is $O(T) + O(1/\varepsilon)$. Let $\varepsilon = (1-\rho)/2$. Then

$$\frac{1}{\varepsilon} = \frac{1}{1-(1+\varepsilon)\rho} \leq \frac{2}{1-\rho}. \tag{21}$$

That is, this choice of ε increases the upper bound on the delay performance by only a constant factor with respect to the desired lower bound. The existence of such a constant satisfies the order condition and Theorem 3 is proven.

4.3. Comparison of Performance to Lower Bound

First we note that the batching policy above meets the criteria for Theorem 1 to apply. That is, $E[I_k|B_k] = c_1\sqrt{A}\sqrt{B_k} + 2\sqrt{2}\sqrt{A}/v$ which satisfies condition **(c)** of the Corollary.

Note that for $r = 0$, the delay of the policy above approaches that of the lower bound (up to a constant) arbitrarily closely as $\varepsilon \to 0$. However, letting $\varepsilon = (1-\rho)/2$ where ρ is close to 1,

$$\frac{1}{\varepsilon} \approx \frac{1}{1-(1+\varepsilon)\rho} \approx \frac{2}{1-\rho}. \tag{22}$$

That is, this choice of ε increases the upper bound on the delay performance by only a constant factor with respect to the desired lower bound. Therefore, we can say that this bound is tight up to a constant when $\rho \to 1$.

5. Discussion

Below we briefly comments on a few extensions of the policy analysis and the lower bounds.

5.1. The 1-Depot DPDP

In addition to the synchronous policy described above, we may also consider a policy in which messages may be dropped at a depot for asynchronous relay service. For this, we assume the existence of a fixed-location depot at which each vehicle may drop an unlimited number of messages at any time. Messages remain at the depot until they are picked up at the depot for delivery by other vehicles. It is assumed that multiple vehicles may transmit simultaneously to and from the depot while at the depot location.

The assignment and service components of the 1-Depot policy are similar to those of the synchronous policy in Section 4 with $r=0$ which we have already seen to be order optimal. Allowing asynchrous relays implies that variable size batches may be used and $a-$optimal TSP tours may be used to visit the service locations instead of worst-case. The $a-$optimal Euclidean TSP tours may be computed using the PTAS in Arora (1998).

The asynchronous nature of the depot relay makes the analysis of this policy more complicated, but the delay of this policy can only be smaller than that of the synchronous policy in which worst case batching and tours are used. In practice, if a depot is available, it can reduce the delay below that of the synchronous policy, although both are of the same order as the lower bound.

5.2. The No-Relay DPDP

According to our previous work on the DPDP with no relays, at least one relay is required to meet the DTRP lower bound. In Waisanen et al. (2006), the average delay for the DPDP without relays scales as $\Theta(\frac{\lambda(n)A}{v^2(1-\rho)^2 n^{3/2}})$ with $\rho = 2\lambda(n)\bar{s}(n)/n$. This is a factor of $1/\sqrt{n}$ weaker than the relay results presented here.

The main difference between the relay policies and the no relay policies is in the assignment subregion served by each vehicle. When no relays are allowed, each vehicle must serve two regions: pickup and delivery. Because messages arriving to a single subregion are destined throughout the entire region, vehicle pickup regions must overlap. This implies that each vehicle must cover a region of area at least $A/\sqrt{n}$. In the relay policy, pickup regions need not overlap, and each vehicle services only a region of area A/n, decreasing the area that must be covered in the pickup and delivery tours, thus decreasing delay.

However note that if ρ is large ($\rho \geq 1/2$) in the DPDP with no relays, the addition of even a single relay will make the system unstable due to the additional onsite service times induced by the relay and the corresponding change in effective ρ. Therefore, in heavily loaded systems, the DTRP bound may not be achievable.

5.3. Sufficiency of Single Relay

To show that a single relay is sufficient to achieve optimal delay, we consider the case in which multiple relays are allowed. The lower bound of the multiple relay problem is the same as for the single relay problem aside from a possible increase in ρ. The Universal Lower Bound is governed only by the distance between source and destinations of individual messages, regardless of the service structure, and the DTRP lower bound already assumes a single service. Because this lower bound is achieved with a single relay, no order improvements are possible with multiple relays. In a multiple-relay system, any savings in relay travel time must be lost when each message is required to go through multiple relays to reach its destination.

5.4. Extension to Multiple Stages

Under certain natural assumptions, the bound in Theorem 2 would hold for the general multi-stage system with two differences. First, if μ services are required per demand, ρ would increase to $\rho = \mu\lambda(n)\bar{s}(n)/n$. We have already seen this increase in ρ seen in the progression from DTRP to DPDP to single relay DPDP. Second, the numerator in equation (2) would be replaced by the appropriate minimum travel time per message. This may be bounded by computing a minimum expected TSP time through the service locations of a single demand.

Therefore, with minor changes, the lower bound methods presented in this paper may be adapted for other multi-stage problems. Similar batching policies may also be implemented, although the nature and performance of these policies naturally depend on the application involved.

Appendix A: Proof of Little's Law for the Onsite System

The onsite system for each vehicle is defined to be the system formed by deleting all time in which the vehicle is not in onsite service. Messages that arrive while the vehicle is in service arrive immediately to the onsite system. Messages that arrive while the vehicle is not in service arrive to the onsite system as soon as a new message begins service. Recall that each message requires two services by the vehicle, and so messages depart the system when they receive their second service by the vehicle. The following relation between the average number in the onsite system at vehicle i, $N_{O,i}$ and the average waiting time in this system, $W_{O,i}$, is used in the proof of Lemma 1.

LEMMA 5 **(Little's Law for the Onsite System.)**. *For the reduced onsite system, the following relation holds on a sample path basis with probability 1:*

$$N_{O,i} = \frac{W_{O,i}}{2\bar{s}(n)}$$

The proof of the lemma is adapted directly from the proof of the sample path version of Little's Law found in Wolff (1989) (pp.286-8). The main change made here is to state Little's Law in terms of departure rate instead of arrival rate. In this appendix, the reference to the vehicle index i is dropped. Further, we will use the time index τ instead of t to denote time in the onsite system.

Recall the definition of the onsite system time for a stable system:

$$W_O = \lim_{J\to\infty} \frac{\sum_{j=1}^{J} W_O(j)}{J} < \infty. \tag{23}$$

$N_O(\tau)$ is defined to be the number in the onsite system at time τ. We assume that the time average number in the system while the system is in service, N_O, is also well defined as

$$N_O = \lim_{\tau\to\infty} \frac{\int_0^{\tau} N_O(\zeta)d\zeta}{\tau}. \tag{24}$$

Further, we show that the following pointwise limit exists with probability 1:

LEMMA 6.

$$\lim_{\tau\to\infty} \frac{N_O(\tau)}{\tau} = 0.$$

Lemma 6 Suppose not. Then for any $\varepsilon > 0$, there exists a τ_ε such that $N_O(\tau) > \varepsilon\tau$ for all $\tau \geq \tau_\varepsilon$. Fix some $\varepsilon > 0$. Let $(\tau_m)_{m=1}^{\infty}$ be a strictly increasing sequence of epochs such that $N_O(\tau_m) > \varepsilon\tau_m, \forall m$. At most one message may leave the onsite system during the interval $[\tau_m, \tau_m + \bar{s}]$, therefore, $N_O(\tau) > \varepsilon\tau_m - 1$ over the entire interval. That is, associated with each τ_m we have the following integral:

$$\frac{\int_{\tau_m}^{\tau_m+\bar{s}} N_O(\zeta)d\zeta}{\tau_m} > \frac{\bar{s}(\varepsilon\tau_m - 1)}{\tau_m} = \bar{s}\varepsilon - \frac{1}{\tau_m}$$

Because $\tau_m \to \infty$ as $m \to \infty$, the limit (24) is equivalent to

$$\begin{aligned}\lim_{m\to\infty} \frac{\int_0^{\tau_m+\bar{s}} N_O(\zeta)d\zeta}{\tau_m + \bar{s}} &= \lim_{m\to\infty} \frac{\int_0^{\tau_m} N_O(\zeta)d\zeta}{\tau_m} \frac{\tau_m}{\tau_m + \bar{s}} + \frac{\int_{\tau_m}^{\tau_m+\bar{s}} N_O(\zeta)d\zeta}{\tau_m} \frac{\tau_m}{\tau_m + \bar{s}} \\ N_O &= N_O + \bar{s}\varepsilon\end{aligned}$$

Since this is true for any $\varepsilon > 0$, this is a contradiction, and (24) is proven.

Let τ_j be the arrival time of the j^{th} message to the onsite system. The departure time of this message is then $\tau_j + W_O(j)$. Define the departure process, $D(\tau)$, to be the number of messages that have departed the onsite system in the interval $[0,\tau]$, that is, $D(\tau) = card(\{j \mid \tau_j + W_O(j) \leq \tau\})$. Let $R(\tau)$ be the number of messages that have completed one full vehicle service, but have not yet left the system (recall that each message must receive two services by the vehicle).

Because the vehicle is continuously serving messages, each requiring deterministic service time $\bar{s}(n)$, the total service rate is given by

$$\lim_{t\to\infty} \frac{D(t)+R(t)}{t} = \frac{1}{\bar{s}(n)}.$$

Because every message must receive at least one service before departing, $D(t) \leq R(t), \forall t$. Further, because the messages in the onsite system may include messages that have not yet received any service, $R(t) \leq D(t) + N_O(t), \forall t$. Combining these and taking limits yields

$$\limsup_{\tau\to\infty} \frac{2D(\tau)}{\tau} \leq \lim_{\tau\to\infty} \frac{D(\tau)+R(\tau)}{\tau} \leq \liminf_{\tau\to\infty} \frac{2D(\tau)+N_O(\tau)}{\tau}$$

From Lemma 6, $N_O(\tau)/\tau \to 0$, and the right-hand and left-hand limits imply that

$$\lim_{\tau\to\infty} \frac{D(\tau)}{\tau} = \frac{1}{2\bar{s}(n)}. \tag{25}$$

With the departure rate precisely established, we move to computing bounds on N_O.

Bound $N_O = \int_0^t N_O(\tau)d\tau$ in the following way:

$$\sum_{j=1}^{D(\tau)} W_O(j) \leq \int_0^\tau N_O(\zeta)d\zeta \leq \sum_{\tau_j \leq \tau} W_O(j). \tag{26}$$

The left-hand inequality counts those customers who have already departed the system whereas the right-hand inequality also includes those who have arrived by time t but have not yet departed. Dividing the middle term by τ and taking the limit as $\tau \to \infty$ yields N_O.

First look at the left hand inequality. Combining (23) and the fact that $D(\tau) \to \infty$ as $\tau \to \infty$ implies that

$$\lim_{\tau\to\infty} \frac{\sum_{j=1}^{D(\tau)} W_O(j)}{D(\tau)} = W_O$$

With this and (25), we divide by τ, take the limit and derive the following result:

$$\lim_{\tau\to\infty} \frac{\sum_{j=1}^{D(\tau)} W_O(j)}{\tau} = \lim_{\tau\to\infty} \frac{D(\tau)}{\tau} \frac{\sum_{j=1}^{D(\tau)} W_O(j)}{D(\tau)} = \frac{W_O}{2\bar{s}}$$

If we can show that a similar bound holds for the right hand inequality, the proof is complete. Again, the existence of the limit in (23) implies

$$\lim_{J\to\infty} \frac{W_O(J)}{J} = \lim_{J\to\infty} \left[\frac{\sum_{j=1}^{J} W_O(j)}{J} - \left(\frac{\sum_{j=1}^{J-1} W_O(j)}{J-1} \right) \left(\frac{J-1}{J} \right) \right] = 0 \tag{27}$$

The Lemma 6 on the number in system process implies

$$\lim_{J\to\infty} \frac{J}{\tau_J} \leq \lim_{J\to\infty} \frac{D(\tau_J)}{\tau_J} + \frac{N(\tau_j)}{\tau_J} = \frac{1}{2\bar{s}(n)} \tag{28}$$

Combining (27) and (28) yields the following limit:

$$\lim_{J\to\infty}\frac{W_O(J)}{\tau_J}=\lim_{J\to\infty}\frac{J}{\tau_J}\frac{W_O(j)}{J}=0$$

This limit implies that for any $\varepsilon>0$, there exists a K such that $W_O(j)\le\tau_j\varepsilon$ for all $j>K$. That is, a message arriving at time τ_j will have departed by time $\tau_j+W_j\le(1+\varepsilon)\tau_j$. We may use this to further upper bound the right hand inequality in (26).

$$\sum_{\tau_j\le t}W_O(j)\le\sum_{j=K+1}^{D(\tau(1+\varepsilon))}W_O(j)+\sum_{j=1}^{K}W_O(j)$$

Dividing by τ and taking the limit, this upper bound converges to $\frac{W_O}{2\bar{s}}(1+\varepsilon)$. Since ε may be arbitrarily small, the lemma is proven.

Appendix B: Proof of Corollary 1

[Corollary 1] If condition **(a)** holds, then $E[B_kI_k]\le E[B_k]E[I_k]$. Then equation (8) is equivalent to

$$\frac{E[B_k^2]}{E[B_k]}\ge\frac{E[B_k]E[I_k]}{E[I_k]}\Leftrightarrow E[B_k^2]\ge E[B_k]^2 \tag{29}$$

since $E[X^2]\ge E[X]^2$ for any random variable X.

If condition **(b)** holds, then equation (8) is equivalent to either

$$\frac{N_B^2}{N_B}\ge\frac{N_BE[I_k]}{E[I_k]}\Leftrightarrow N_B\ge N_B \tag{30}$$

or

$$\frac{\gamma^2E[I_k^2]}{\gamma E[I_k]}\ge\frac{\gamma E[I_k^2]}{E[I_k]}. \tag{31}$$

To show the proof when condition **(c)** holds, we first have the following set of equalities implied by the bounds on $E[I_k|B_k]$:

$$\begin{aligned}E[B_k^2]E[I_k]&=E[B_k^2]E[E[I_k|B_k]]=\gamma E[B_k^2]E[B_k^\alpha]+E[B_k^2]\delta\\ E[B_kI_k]E[B_k]&=E[E[B_kI_k|B_k]]E[B_k]=\gamma E[B_k^{1+\alpha}]E[B_k]+E[B_k]^2\delta\end{aligned}$$

Again, because $E[B_k^2]\ge E[B_k]^2$, the δ terms cancel and it remains to show

$$E[B_k^2]E[B_k^\alpha]\ge E[B_k^{1+\alpha}]E[B_k] \tag{32}$$

for $0\le\alpha\le1$.

Let $p_q=P(B_k=q)$. Then, expanding (32) in terms of the p_q yields

$$\begin{aligned}E[B_k^2]E[B_k^\alpha]&=\left(\sum_{q=1}^{\infty}q^2p_q\right)\left(\sum_{q=1}^{\infty}q^\alpha p_q\right)\\&=\sum_{q=1}^{\infty}\sum_{r=1}^{\infty}q^2r^\alpha p_qp_r\\ E[B_k^{1+\alpha}]E[B_k]&=\left(\sum_{q=1}^{\infty}q^{1+\alpha}p_q\right)\left(\sum_{q=1}^{\infty}qp_q\right)\\&=\sum_{q=1}^{\infty}\sum_{r=1}^{\infty}q^{1+\alpha}rp_qp_r\end{aligned}$$

Matching the terms multiplied by $p_q p_r = p_r p_q$, it remains to show

$$q^2 r^\alpha + r^2 q^\alpha \geq q^{1+\alpha} r + r^{1+\alpha} q \tag{33}$$

for all pairs $(q, r) \in \mathbb{Z}^+ \times \mathbb{Z}^+$. Without loss of generality, assume $q \geq r$. For $\alpha \in [0, 1]$, we have the following series of equivalent inequalities:

$$\begin{aligned}
q\left(q^{1-\alpha} - r^{1-\alpha}\right) &\geq r\left(q^{1-\alpha} - r^{1-\alpha}\right) \\
q^{2-\alpha} - qr^{1-\alpha} &\geq rq^{1-\alpha} - r^{2-\alpha} \\
q^{2-\alpha} + r^{2-\alpha} &\geq rq^{1-\alpha} + qr^{1-\alpha} \\
q^\alpha r^\alpha \left(q^{2-\alpha} + r^{2-\alpha}\right) &\geq q^\alpha r^\alpha \left(rq^{1-\alpha} + qr^{1-\alpha}\right) \\
q^2 r^\alpha + r^2 q^\alpha &\geq q^{1+\alpha} r + r^{1+\alpha} q
\end{aligned}$$

Therefore, the theorem holds when condition **(c)** is true.

Remark: Note that the corollary holds when I_k is affine in B_k. However this does not hold if B_k is affine in I_k. In that case

$$\begin{aligned}
E[(aI_k + b)^2]E[I_k] &= a^2 E[I_k^2]E[I_k] + 2abE[I_k]^2 + b^2 E[I_k] \\
E[(aI_k + b)I_k](aE[I_k] + b) &= a^2 E[I_k^2]E[I_k] + abE[I_k]^2 + abE[I_k^2] + b^2 E[I_k] \\
\text{but} \quad abE[I_k]^2 &\leq abE[I_k^2].
\end{aligned}$$

Acknowledgments

This material is based upon work supported under a National Science Foundation Graduate Research Fellowship as well as a National Defense Science and Engineering Graduate Fellowship.

References

Arora, S. 1998. Polynomial-time approximation schemes for euclidean tsp and other geometric problems. *Journal of the ACM* **45**(5) 753–782.

Bertsimas, D.J., G. van Ryzin. 1991. A stochastic and dynamic vehicle routing problem in the euclidean place. *Oper Res* **39**(4) 601–615.

Bertsimas, D.J., G. van Ryzin. 1993a. Stochastic and dynamic vehicle routing in the euclidean plane: the multiple-server, capacitated vehicle case. *Oper Res* **41** 60–76.

Bertsimas, D.J., G. van Ryzin. 1993b. Stochastic and dynamic vehicle routing with general demand and interarrival time distributions. *Adv App Prob* **20**(4) 947–978.

Desaulniers, G., J. Desrosiers, A. Erdmann, M.M. Solomon, F. Soumis. 2002. Vrp with pickup and delivery. P. Toth, D. Vigo, eds., *The Vehicle Routing Problem*. SIAM Monographs on Discrete Mathematics and Applications, 225–242.

Feuerstein, E., L. Stougie. 2001. On-line single-server dial-a-ride problems. *Theo Comp Sci* **268** 91–105.

Frazzoli, E., F. Bullo. 2004. Decentralized algorithms for vehicle routing in a stochastic time-varying environment. *IEEE Conf. on Decision and Control*. Paradise Island, Bahamas, 3357–3363.

Gendreau, M., J-Y Potvin. 1998. Dynamic vehicle routing and dispatching. T.G. Crainic, G. Laporte, eds., *Fleet Management and Logistics*. Kluwer: Boston, 115–126.

Larson, R., A. Odoni. 1981. *Urban Operations Research*. Prentice-Hall, Englewood Cliffs, NJ.

Savelsbergh, M.W.P, M. Sol. 1995. The general pickup and delivery problem. *Transportation Science* **29**(1) 17–29.

Savla, K., F. Bullo, E. Frazzoli. 2006. Traveling Salesperson Problems for a double integrator. *IEEE Transmissions on Automatic Control* Submitted.

Swihart, M.R., J.D. Papastavrou. 1999. A stochastic and dynamic model for the single-vehicle pick-up and delivery problem. *Eur J Oper Res* **114** 447–464.

Waisanen, H. A., D. Shah, M. A. Dahleh. 2006. A dynamic pickup and delivery problem in mobile networks under information constraints. *IEEE Transactions on Automatic Control* Submitted.

Wolff, R.W. 1989. *Stochastic Modeling and the Theory of Queues*. Prentice-Hall, Englewood Cliffs, NJ.